LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

LIBRA

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

MAIL-ORDER DEMANDS FOR STYLE GOODS:
THEORY AND DATA ANALYSIS

W.H. Hausman and R. St.G. Sides

611-72 July, 1972

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

DEWEY LIBRARY

MAIL-ORDER DEMANDS FOR STYLE GOODS:
THEORY AND DATA ANALYSIS

W.H. Hausman and R. St.G. Sides

611-72 July, 1972

HD28
.M414
no. 615-72

Dewey

MASS. INST. TECH.
SEP 23 1975
LIBRARIES

RECEIVED
OCT 17 1972
M. I. T. LIBRARIES

75 333HY 13 061 BB 4535

ABSTRACT

A number of alternative data-generating processes are explored for mail-order demands for seasonal style-goods. Weekly demands for 126 items over an 18-week selling season provide the empirical data.

Arguments are presented which result in the following candidates for data-generating processes: (1) ratios of successive forecasts are distributed lognormally; (2) ratios of successive forecasts are distributed as _t_ (Student); (3) actual demands during unequal time periods are distributed as negative binomial. Analysis of the data suggests the negative binomial data-generating process as both most closely representing the underlying process and being simple to adapt to a decision model. The paper concludes with an example of the use of the chosen data-generating process in a decision framework, and deals briefly with some issues of implementation.

636041

I. Introduction

We consider seasonal style goods with a limited selling season which are ordered by mail from a catalog. Our goal is to achieve a mathematical model representing the demand data-generating process for such items over the selling season. This model can be subsequently combined with a decision framework dealing with when and how much to order of each item, given various costs of oversupply and stockout.

Management of seasonal style-goods inventory has been studied by a number of authors (see [2], [3], [6], [7], [8], [9], [12], [16]). A review of most of this literature is given in [6] and will not be repeated, except to summarize that much of the previous work cited presents models without testing the assumptions concerning the data-generating process for demands.[1] References [8] and [16] do deal with actual data: demands for women's apparel at the retail store level. These studies reveal that retail sales tend to be proportional to the amount of inventory displayed; however, the mail-order case never runs out of display inventory, and is thus fundamentally different from the retail case. Reference [7] deals with data generalized from wholesale demands placed on a textile manufacturer, which again is quite distinct from the retail mail-order case.

The plan of the paper is as follows. Sections II, III and IV will describe three alternative models for the demand data-generating process, with empirical results suggesting the third model as being most appropriate. Section V combines the chosen model with an economic decision framework for when and how much to order, and deals briefly with some issues of implementation.

---

[1] In particular, [12] assumes a continually increasing mean demand rate throughout the entire season, an assumption flatly contradicted by the empirical results of [7], [8], and [16].

II. The Lognormal Model

Instead of studying demands for seasonal goods directly, we may focus on statistical properties of forecasts (and forecast revisions) of total seasonal demand for each item, where the forecast mechanism is either a simple extrapolation of sales-to-date (as in [5] and [7]) or a more complex Bayesian combination of an initial, prior forecast and experience to date (as in [2], [3] and [9]). Previous empirical work [5] has indicated that under certain circumstances, systems for generating and revising forecasts of some unknown quantity may have the property that ratios of successive forecasts can, as a first approximation, be treated as independent random variables distributed according to the Lognormal distribution.

Specifically, we divide the season into periods $j = 1,\dots,N$. Let X_j represent the forecast of total seasonal demand for a good, with the forecast made at the end of period j; also let X_N represent actual total demand for the good. Define

$$(1)\qquad Z_j = X_{j+1}/X_j \ , \ j = 1,\dots,N-1 \ .$$

Then the above assumptions may be restated as follows; Z_j has the distribution

$$(2)\qquad Z_j \simeq f_{LN}(Z_j/\mu_j,\sigma_j) = \frac{1}{(2\pi)^{1/2}\sigma_j Z_j}\ e^{-(\log Z_j-\mu_j)^2/2\sigma_j^2} \ , \ j = 1,\dots,N-1$$

with Z_j independent of Z_k for $j \neq k$.

The Data

Before the above assumptions are tested, we briefly describe the nature of the data. Weekly mail-order demands for 126 separate items of women's sportswear were obtained over an 18-week selling season. The

aggregate demand pattern (across items) over the season is graphed in Figure 1.

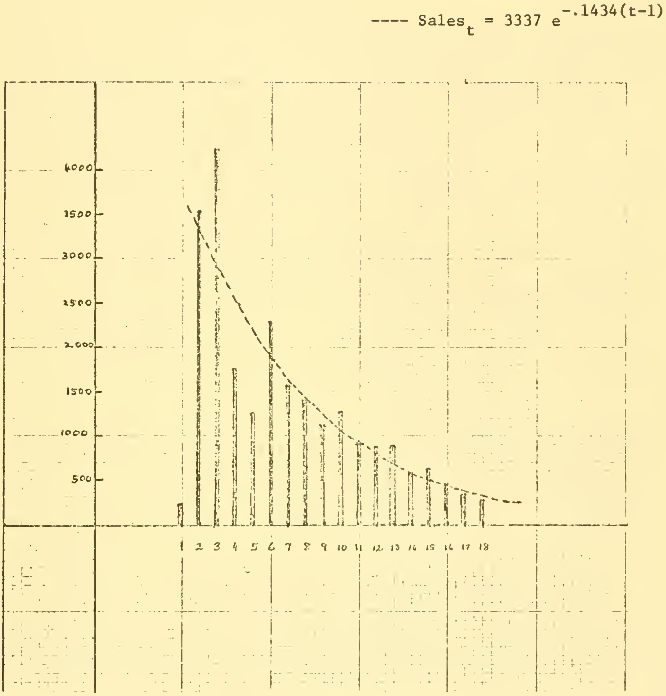

Figure 1: Aggregate Mail-Order Demand by Week

Since aggregate demand in week one and in the latter weeks of the season was quite low, the weeks were consolidated into N = 9 different (unequal) time periods in an attempt to gain rough uniformity of the aggregate demand rate per time period.

The lognormality assumption was tested against forecasts of total seasonal demand for each item, with each forecast being revised after each time period. The forecasting method was a simple extrapolation of item demands to date, using the aggregate percentage of demands in each of the nine periods as the basis for extrapolation.[2] Then ratios Z_j of successive forecasts of total demand for each item were computed as in equation (1), and the natural logarithm $W_j = \ln(Z_j)$ was computed for all items. If Z_j is distributed Lognormally, then by definition its natural logarithm is distributed Normally:

$$(3) \qquad W_j \simeq f_N(W_j | \mu_j, \sigma_j) = \frac{1}{(2\pi)^{1/2}\sigma_j} e^{-(W_j-\mu_j)^2/2\sigma_j^2} \quad , \; j = 1,\ldots,N-1 \; .$$

Figures 2, 3 and 4 illustrate ordered plots of W_1, W_2 and W_8 on Normal probability paper. The probability paper is constructed such that a true Normal variate would lie on a straight line except for sampling variation.

A reasonable fit occurs for $W_8 = \ln(Z_8) = \ln(X_9/X_8)$ but the other examples, W_1 and W_2, show an unsatisfactory curve with a very high right tail. (Sample third and fourth moments also indicate significant departures from Normality.)

---

[2] There is a minor amount of circular computing here, but we felt that 126 items was a sufficient number so that the effect of demands of any one item on the aggregate percentages could be ignored.

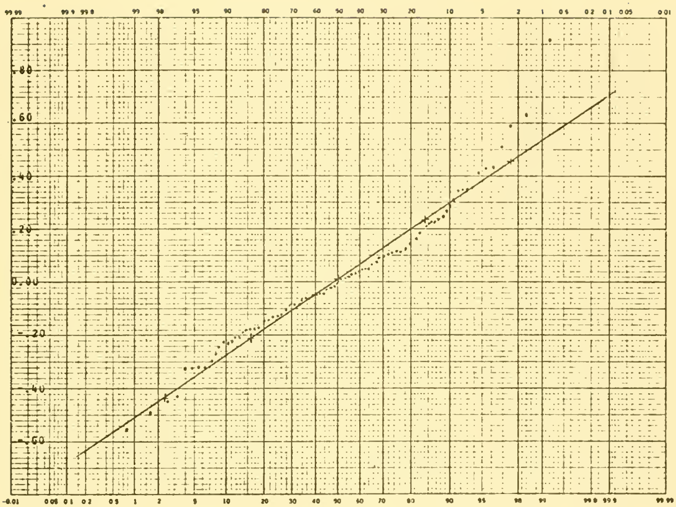

Figure 2: W_1 = log of X_2/X_1

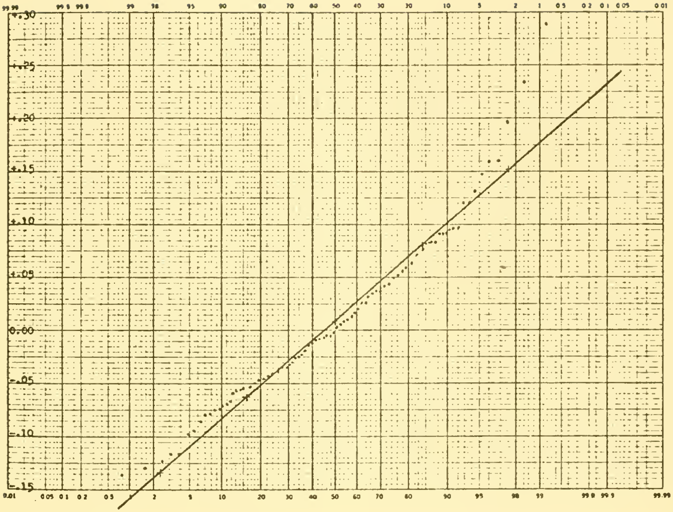

Figure 3: W_2 = log of X_3/X_2

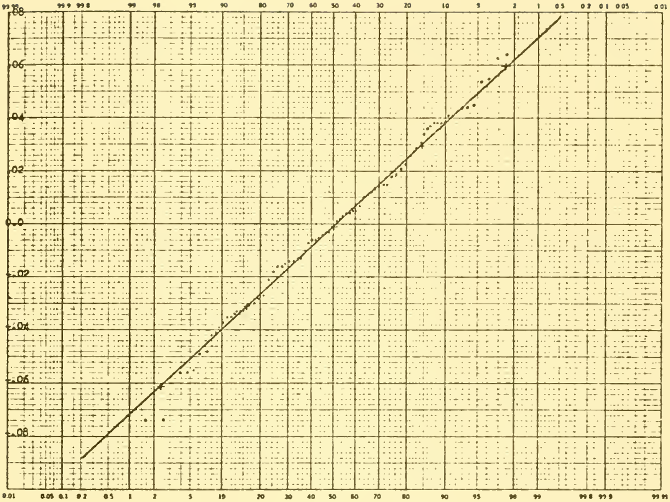

Figure 4: W_8 = log of X_9/X_8

The Lognormality assumption implies equal "percentage variability" across items selling at substantially different rates. This aspect of the Lognormality assumption was tested by treating separately 51 low-unit volume items and 51 high-unit volume items (omitting 24 items with pre-season sales from a previous catalog). Table 1 contains the ratio of the standard deviation of W_j for the low-volume items to the standard deviation of W_j for the high-volume items.

Table 1

$$\text{Ratio} = \frac{\text{Standard Deviation of } W_j \text{ for 51 low-volume items}}{\text{Standard Deviation of } W_j \text{ for 51 high-volume items}}$$

j:	1	2	3	4	5	6	7	8	Mean
Ratio:	2.109	1.773	2.120	1.415	2.181	1.904	1.642	1.762	1.863

The relative variability in the forecast revision process for low-volume items is substantially larger than that for the high-volume items, a phenomenon frequently found in demand analyses. Treating all 126 items as being from a single population thus discards information that relative variability of forecast revisions is inversely related to the average demand rate for an item.

The independence assumption (that Z_j is independent of Z_k for $j \neq k$) was tested by performing a correlation analysis on the W_j's (see Table 2).

Table 2

Intercorrelation Matrix for Logs of Ratios of Successive Forecasts

	W_1	W_2	W_3	W_4	W_5	W_6	W_7
W_1	1.0000						
W_2	-0.0811	1.0000					
W_3	-0.1214	0.0705	1.0000				
W_4	0.0596	-0.2439	-0.0692	1.0000			
W_5	-0.0662	-0.0943	-0.0012	0.1527	1.0000		
W_6	0.1016	-0.1132	-0.0395	0.1047	0.1529	1.0000	
W_7	0.0928	0.1062	0.0601	-0.1984	0.1537	0.0391	1.0000
W_8	0.0058	-0.1615	-0.0448	0.1468	0.4012	0.2728	0.1094

Table 2 contains three correlations greater than 0.230 in absolute value, which represents the 2% significance level for $n = 126$. This would indicate the hypothesis of mutual independence of forecast revisions across time periods to be false; however, no causal explanation has been found for the particular non-zero correlations (r_{24}, r_{58}, r_{68}). One must have continuing systematic departure from independence in order to be able to make use of the dependency. We therefore interpret the test of independence as being inconclusive.

Based on Figures 2 through 4 and Table 1's results, the Independent Lognormal model is not adequately supported by this data.

III. The Heterogeneous Lognormal Model (*t* - distribution)

As noted above, the standard deviation σ_j of each of the W_j's is not constant across items. Suppose that parameter were considered as a random variable. Then, having chosen the form and parameters of the distribution of σ_j, that distribution may be combined with the (conditional) Lognormal distribution above, and integration over the range of values for σ_j would produce the marginal distribution of W_j. This precise approach was recently used by Praetz [10] in connection with the distribution of stock price changes.

The chosen distribution for σ_j must be such that one can perform the indicated integration, and for this reason we use the corresponding Bayesian natural conjugate prior distribution (see [11]). However, this borrowing of a mathematically tractable distribution from Bayesian decision theory does not mean that a Bayesian approach is being used in this model. To the contrary, we (as Praetz) argue in this section that there is underlying variability in the parameter in question (σ_j), which is not reduced by any sample information but which simply remains throughout the process.

Specifically, equation (3) is now interpreted to be the conditional probability density function for W_j, conditional on σ_j;

$$(4) \qquad f_N(W_j|\sigma_j) = \frac{1}{(2\pi)^{1/2}\sigma_j} e^{-(W_j-\mu_j)^2/2\sigma_j^2} \quad .$$

As a conjugate prior distribution we follow Praetz [10] and assign an inverted-gamma-one[3] distribution on σ_j^2:

---

[3] See [10], p. 51, and [11], p. 227.

$$(5) \qquad f_{i\gamma1}(\sigma_j^{\,2}|m,(m-1)\sigma_0^{\,2}) = \sigma_0^{\,2}(m-1)^m\sigma^{-2(m+1)}e^{-(m-1)\sigma_0^{\,2}/\sigma^2}\,/\Gamma(m)$$

with $E[\sigma^2] = \sigma_0^{\,2}$ and Var $[\sigma^2] = \sigma_0^{\,4}/(m-2)$. Then the unconditional distribution of W_j is, following [10], a Student distribution:

$$(6) \qquad f_S(W_j) = \int_0^\infty f_N(W_j|\sigma_j)f_{i\gamma1}(\sigma_j^{\,2}|m,(m-1)\sigma_0^{\,2})\,d\sigma^2$$

$$= [1+(W_j-\mu_j)^2/\sigma_0^{\,2}(2m-2)]^{-m-\frac{1}{2}}\;\Gamma(m)[(2m-2)\pi]^{\frac{1}{2}}\;\sigma_0 :$$

> "... this is a *t* distribution of $2m = n$ degrees of freedom, except for a scale factor $[n/(n-2)]^{1/2}$. Thus, the distribution of $[(W_j-\mu_j)/\sigma_0]$ would be that of a scaled *t*-distribution. When this distribution is plotted against a standard normal distribution for *n* small, it reproduces the characteristic distribution of share price changes described earlier" [10, pp. 50-51].

Figures 5 and 6 contain plots of scaled *t*-distributions with $n = 4$ and $n = 6$ degrees of freedom on Normal probability paper. Comparison with Figures 2 through 4 indicates that the shape of the *t*-distribution would generally provide a better fit to the W_j data than the Normal distribution.[4]

Praetz fitted *t*-distributions to logarithms of share price ratios for seventeen indices, using weekly data from the Sydney Stock Exchange. Full details are in [10] but for the seventeen series, Praetz found the best *n* value (degrees of freedom) was 3 twice, 4 twelve times, 5 twice and 7 once. Thus, for logarithms of ratios of successive index values for stocks, a

---

[4] Comparison of Figures 5 and 6 with Figures 8 through 10 of [5] also indicate that the *t*-distribution would tend to fit those sample W_j values better than the Normal distribution.

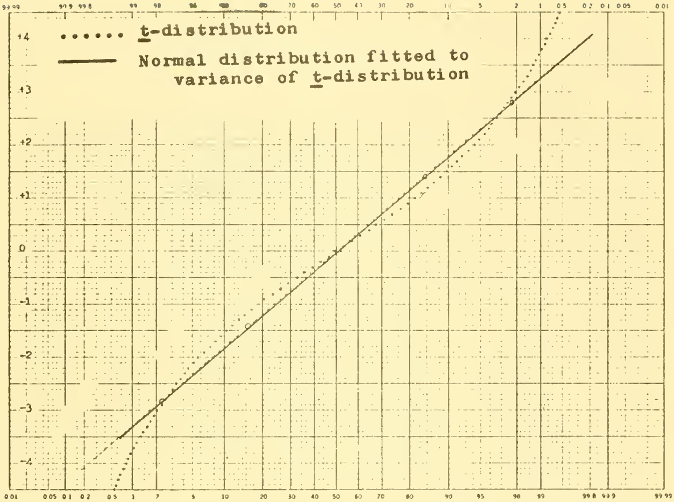

Figure 5: t-distribution with n = 4 degrees of freedom

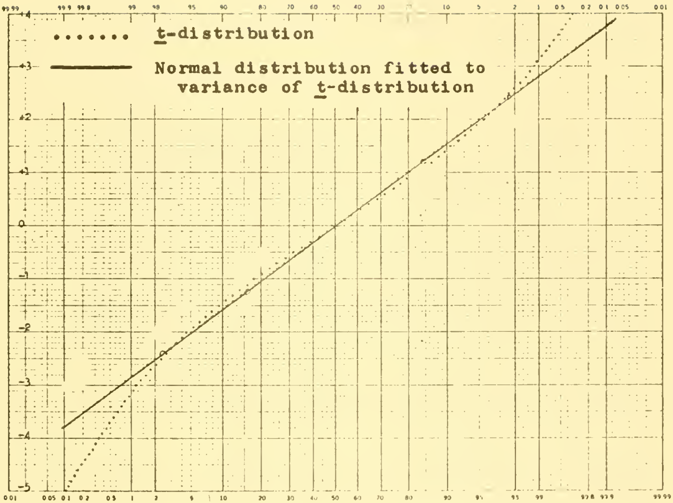

Figure 6: t-distribution with n = 6 degrees of freedom

t-distribution with approximately n = 4 degrees of freedom seems to fit the data quite well, substantially better than other candidates such as the Normal distribution, the Stable Paretian distribution, and the Compound Events distribution [10, p. 53].

With this background, the t-distribution was fitted to our W_j data; see Appendix 1 for the details. Table 3 contains the degrees-of-freedom value which provided the best fit in each case.

Table 3

Degrees of Freedom for Fitted t-distribution

Data:	W_1	W_2	W_3	W_4	W_5	W_6	W_7	W_8
Degrees of Freedom (n):	4	11	3 or 4	6	3	5	6	6

Thus small values of n generally provide the best fits for our data, as was found in [10].[5]

As compared to the Independent Lognormal Model, this Heterogeneous Lognormal Model leading to a t-distribution of W_j fits the data better. However, in contrast to Praetz, we have already found a reason for the apparent variability in σ across the 126 items in our sample; namely, that relative variability is inversely related to the average demand rate for the item. For this reason, although the second model is an improvement

---

[5] Values of n less than 5 may appear to create difficulties with the prior distribution of σ^2 in equation (5), since n=4 implies m=2 and the variance of σ^2 is then undefined. The solution is to use $h = 1/\sigma^2$ as a transformed variable, assessing a gamma-two distribution on h [11, pp. 226,291] to avoid this difficulty.

over the first, we turn to a third model capable of incorporating the relationship of relative variability to average demand rate.

IV. The Heterogeneous Poisson Model (Negative-Binomial Distribution)

Instead of considering the statistical behavior of forecasts (and reforecasts) of total season demand for each item, we now focus on the number of demands for a given item occurring during a stated number of aggregate transactions[6] (e.g., every 2000 aggregate demands across all items). A simple Poisson model is first explored, followed by a heterogeneous Poisson model.

On a transactions basis the underlying characteristics of the simple Poisson process seem applicable to the mail-order situation; units are ordered (generally) one at a time, and in any given selling season, the average number of demands for a given item in a given time period should be approximately proportional to the total number of transactions in the stated time period. Moreover, since the variance equals the mean in a Poisson process, the measure of relative variability used in the W_j data analysis above (the standard deviation of W_j the logarithm of the ratio of successive forecasts) would theoretically decrease as the average demand rate increased; this does in fact happen, as indicated in Table 1. Thus, the simple Poisson model was a candidate for analysis.

---

[6] We consider transactions rather than calendar time because as Figure 1 indicates, the aggregate demand rate is not constant over the selling season.

Poisson arrivals imply exponential inter-arrival times, as is well-known. If the time between demands for each item were available, this data could be plotted against an exponential distribution and the goodness of fit evaluated statistically. However, inter-arrival data were not available, so less-direct tests were used.

Table 1 showed that the mean value across all periods of the ratio (standard deviation of W_j for 51 low-volume items)/(standard deviation of W_j for 51 high-volume items) was 1.863. The mean demand rate for the 51 high-volume items was 304.1 and for the 51 low-volume items was 87.9. If the standard deviation of the W_j's is considered to represent a measure of variability in underlying demand occurrences, then since the variance equals the mean for a Poisson process, the ratio of standard deviations (1.863) should not be too different from the ratio of the square root of the two mean demand rates: $(304.1/87.9)^{1/2} = (3.459)^{1/2} = 1.860$. This result should be interpreted as being consistent with a Poisson process in that the relationship of means and variances is maintained in a dichotomy grouping.

A second investigation involved considering the nine time periods and 126 items as 1134 realizations. If demand in each period for each item was Poisson, then the expected demand for item i in period j would be estimated by taking the total seasonal demand for item i and pro-rating that by the aggregate percentage of demands occurring in period j. Thus if item 5 had 200 demands in total and if period 6 contained 10% of the total number of orders across all 126 items, then the estimate for expected demand for item 51 in period 6 would be 200(.10) = 20 units.

For each of the 1134 realizations, an expected frequency (e.g., 20) was calculated as indicated. Then, using those 1134 separate values in turn, the probability of 0, 1, 2,...,125 demands occurring was calculated; and these probabilities were accumulated by the number of occurrences. The result was a table of expected frequency of cells containing 0, 1,...,125 demand in the 1134 realizations. Appendix 2 contains a listing of expected and actual frequencies of the various counts over the 1134 realizations; the actual and expected frequencies are quite similar. The cells were consolidated into 27 groups so that a chi-square test could be performed, with a resulting value of chi-square of 12.753; for 24 degrees of freedom[7] this has a P of .975.

The above test is necessary but not sufficient to demonstrate an underlying Poisson process. It is possible for a totally different data-generating process to generate demands in such a way that when analyzed in summary form as above, the results could be identical. Thus, further testing was required.

A third exploration involved computing the ratio of equation (7) for a random sample of 41 items (1/3 of the data):

$$\text{Ratio} = \frac{(\text{observed} - \text{expected})^2}{\text{expected}} \qquad (7)$$

where observed represents the realized item demand of each of the item/period realizations described above, and expected refers to the expected demand

---

[7] One could debate what the appropriate degrees of freedom should be after consolidation, but no matter what the conclusion, this is still a very low value for a chi-square statistic.

calculated in the pro-rata manner described above. Over 41 items and 9 periods (369 observations) the average value for the ratio of equation (7) was 1.194. Now the expected values were distributed over a wide range, averaging 23,591/1134 = 20.8, and the skewness (large or small) of the hypothesized Poisson distributions makes it not possible to make a significance test on the quantity 1.194. One would suspect, however, that since the amount of skewness in a Poisson process with parameter $\lambda = 20.8$ is quite small, the mean ratio value of 1.194 is suspiciously greater than one, thereby indicating that the variance may be greater than the mean.

This leads us to consider the heterogeneous Poisson process in which the Poisson parameter is treated as a random variable with a gamma-one distribution, resulting in the unconditional distribution of demands per period being negative binomial; the latter distribution can represent a situation where the variance/mean ratio is greater than one.

Let us define the following symbols:

t_j = number of aggregate transactions (demands) in period j

D_{ij} = number of demands for item i in period j (a random variable)

K_i = total seasonal demand for item i: $(K_i = \sum_{j=1}^{9} D_{ij})$; (a random variable)

$T_j = \sum_{k=1}^{j} t_k$ cumulative aggregate demands from start of season through period j

$T = T_9$ = total aggregate demands over entire season for all items

= 23,591

λ_i = probability that a transaction is a demand for item i; $\lambda_i = K_i/T$

$S_{ij} = \sum_{k=1}^{j} D_{ik}$ = cumulative demands for item i from season start through period j

$r_{ij} = K_i - S_{ij}$ = demands for remainder of season for item i (a random variable)

$n_j = T - T_j$ = aggregate transactions for remainder of season (presumed known)

Now after a given period j, cumulative demands for item i are given as S_{ij}. We desire the probability function of r_{ij} which represents the remaining demands for item i yet to come in periods j+1,...,N. Presuming the seasonal pattern is reasonably specified,[8] the values of n_j are assumed known. Then, according to the Poisson process assumption, the conditional probability of r_{ij} remaining demands for item i occurring in the next n_j aggregate transactions is:

$$(8) \qquad f_P(r_{ij}|\lambda_i n_j) = e^{-\lambda_i n_j}(\lambda_i n_j)^{r_{ij}}/r_{ij}! \quad .$$

A natural conjugate prior distribution [11, p. 277] for λ_i is the gamma-one distribution

$$(9) \qquad f_{\gamma 1}(\lambda_i|S_{ij},T_j) = \frac{e^{-\lambda_i T_j}(\lambda_i T_j)^{S_{ij}-1}\,T_j}{(S_{ij}-1)!}$$

with $E[\lambda_i] = S_{ij}/T_j$, $\mathrm{Var}(\lambda_i) = S_{ij}/(T_j)^2$, and mode $(\lambda_i) = (S_{ij}-1)/T_j$. This gamma-one distribution of (9) describes our information relevant to λ_i from a sample of aggregate transactions of size T_j containing S_{ij} successes (i.e., demands for item i). Combining (8) and (9), the unconditional distribution of r_{ij} is negative binomial [11, p. 284]:

$$(10) \qquad P\{r_{ij}|S_{ij},T_j,n_j\} = \int_0^\infty f_p(r_{ij}|\lambda_i n_j) f_{\gamma 1}(\lambda_i|S_{ij},T_j)\,d\lambda_i$$

$$= f_{nb}\left(r_{ij}\left|\frac{n_j}{n_j+T_j},\ S_{ij}\right.\right)$$

$$= \frac{(r_{ij}+S_{ij}-1)!}{r_{ij}!(S_{ij}-1)!}\left(\frac{n_j}{n_j+T_j}\right)^{r_{ij}}\left(\frac{T_j}{n_j+T_j}\right)^{S_{ij}}$$

---

[8] Such an assumption is made in [2], [3], [7], and [9].

with $E[r_{ij}] = (n_j S_{ij})/T_j$ and $\mathrm{Var}(r_{ij}) = [(n_j S_{ij})/T_j]\cdot[(n_j+T_j)/T_j]$. The variance of r_{ij} is equal to the mean multiplied by a constant greater than one. In particular, when n_j is large relative to T_j, as at the beginning of the season, this constant is much greater than one, allowing for a large relative initial variance.

Comparison with Beta-Binomial Distribution

The negative-binomial distribution for r_{ij} just derived bears a close relationship to the Beta-binomial distribution used by Murray and Silver [9]. They represent the demand data-generating process by a binomial share-of-market model in which the binomial parameter p is assigned a Beta prior distribution. The resulting unconditional distribution for demands for a particular item during a specified number of transactions is Beta-binomial with parameters:

$$E(r_{ij}) = \frac{n_j S_{ij}}{T_j} \; ; \quad \sigma^2_{r_{ij}} = \frac{n_j S_{ij}}{T_j} \cdot \frac{n_j + T_j}{T_j} \cdot \frac{T_j - S_{ij}}{T_j + 1}$$

The means of the two distributions are the same and the variances are almost identical, differing only by the last term representing the correction for a finite population; in our data the smallest T_j is T_1 = approximately 2000 transactions, and the median S_{i1} is 13.2 so that the value of that last term would be 0.993, a negligible correction. Computationally, the Beta-binomial model requires evaluation of the factorial $(T_j+n_j-1)!$, while the negative-binomial requires $(r_{ij}+S_{ij}-1)!$, a much easier task.

Figure 6 contains a plot of a negative binomial distribution on Normal probability paper for the case when K_i is relatively small (39, or 1/4 of median demands of 156), and for Period 1. Its upward curve is able to capture the more-than-Normal probability of unusually large demands accounting for the skewness present in all the previous figures. The range of this curve to be used in decision-making will generally lie between the 50th and 95th percentiles; note that for the numerical case plotted (the worst for this approximation), the Normal approximation will be too large in most cases but probably not by enough to create serious error. (This approximation is explored further in Section V.)

The use of equation (10) in practice would involve calculating S_{ij} and T_j after each period j, estimating (and perhaps re-estimating) n_j, and using (10) to represent the probability density function of remaining demands for item i during the rest of the selling season.

V. A Decision Model Using the Negative-Binomial Distribution

In theory there may be multiple opportunities to place orders for stock throughout the season, although orders placed later in the season face increasing probabilities of non-fulfillment since manufacturers are generally working on the next season's items by that time. A thorough treatment of the multiple-opportunity-to-buy case would involve a multistage decision framework with solution by dynamic programming as illustrated by [5] and [9].

For the specific circumstances of the mail-order situation used to obtain the data, it is possible to avoid the multistage formulation as

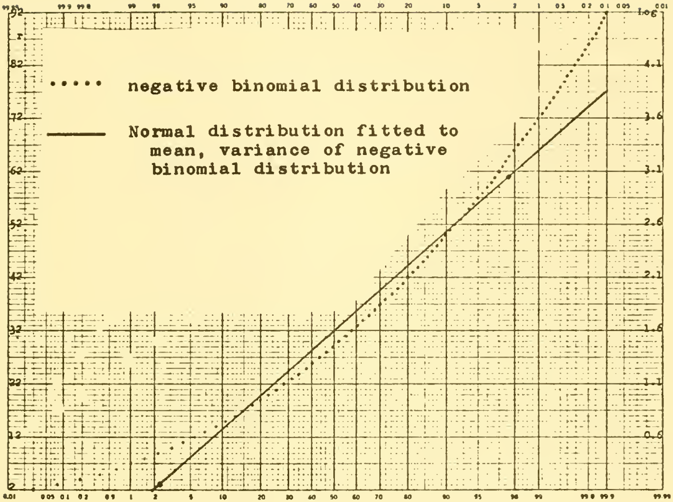

Figure 6: Plot of Negative Binomial Distribution

There is a point in the season beyond which one cannot reasonably expect to obtain delivery of items for this selling season. Delaying the final commitment decreases the cost of uncertainty; against this improvement the buyer must assess the qualitative cost of delaying the final procurement, with its accompanying risks of non-fulfillment, and select (ahead of time) the approximate time in the season for the final procurement. Then that quantity should be determined by the standard newsboy formula, with appropriate parameters; prior commitments can be made essentially according to economic quantity considerations as long as they do not represent an appreciable probability of overstocking for the entire season.

The error involved in using the Normal approximation to the negative-binomial distribution when the latter is the underlying distribution was explored as follows. In Period 1 of our median numerical example, the optimal Q is 173 total units (including 25 used to fill demands in period 1). The normal approximation produces an approximate order quantity of 175 units; the increase in cost of uncertainty resulting from this higher-than-desired inventory level is $0.566, quite a small percentage of the values in Table 4. Thus, if the numbers for this example are at all typical, the

The above methodology is implemented as follows. Each period $j = 1,\ldots,9$ is considered in turn as the last opportunity to order replenishment stock, and the optimal amount to "order up to"

is determined by the usual newsboy formula:

$$F_{nb}\left(Q_j \middle| \frac{n_j}{n_j+T_j}, S_{ij}\right) = \frac{C_u}{C_o+C_u} \tag{11}$$

where C_o is the unit cost of overage, C_u the unit (opportunity) cost of underage, Q_j the optimal inventory level to order up to at the end of period j, and $F_{nb}(\cdot)$ the cumulative negative binomial distribution function of remaining demand over the season for item i. Equation (11) was used to evaluate optimal inventory for the numerical case in which $E[K_i] = 156$ (the median), actual demands D_{ij} follow the pro-rata percent breakdown implied by n_j, $C_u = \$17$, and $C_o = \$7$. The expected cost of uncertainty can be estimated [14, pp. 131-2] by equation (12):

$$\text{Expected Uncertainty Cost} = (C_o+C_u)\ \sigma_{r_{ij}}\ P(r=Q) \quad . \tag{12}$$

Table 4 contains these results.

Table 4

Expected Cost of Uncertainty

Period	1	2	3	4	5	6	7	8	9
Q (whole season)	173	166	163	162	160	160	158	157	156
Cost of Uncertainty	203.59	127.16	109.78	87.25	74.24	57.15	45.22	25.39	0
Marginal Benefit of Delay for One Period	76.43	17.38	22.53	13.01	17.09	11.93	19.83	25.39	--

The benefit relationship would be somewhat smoother if the periods contained identical numbers of transactions; also, because Q is integer the value of P(r=Q) is lumpy.

Table 4 indicates the approximate magnitude of the reduction in the cost of uncertainty as the final commitment is delayed. As stated above, against this improvement the buyer must assesss the qualitative cost of delaying the final procurement, with its accompanying risks of non-fulfillment, and select (ahead of time) the approximate time in the season for the final procurement. Then the optimal order-up-to quantity may be determined by equation (11).

The error involved in using the Normal approximation to the negative-binomial distribution when the latter is the underlying distribution was explored as follows. In Period 1 of our median numerical example, the optimal Q is 173 total units (including 25 used to fill demands in period 1). The normal approximation produces an approximate order quantity of 175 units; the increase in cost of uncertainty resulting from this higher-than-desired inventory level is $0.566, quite a small percentage of the values in Table 4. Thus, if the numbers for this example are at all typical, the

statistical derivations and tests notwithstanding, a simple normal approximation produces decision values close enough to the optimal ones to ignore the difference.

Issues of Implementation

The data analyzed above represented demand by item; however, the items (women's sportswear) come in sizes and colors. It is the custom of the trade to forecast by item and then to disaggregate the order into sizes and colors using various *ad hoc* rules. The argument given for choosing the item as the unit for analysis is that the distribution of sizes is relatively constant and that sales of individual stock-keeping units (item by size and color) are highly correlated. Conversely, further aggregation of items into classes of garment (for example) would not contribute any useful information since demands across items tend to be quite variable. The latter point was explored in some detail by correlating the values W_j (the logarithms of ratios of successive forecasts of total seasonal demand by item) across *items* rather than across time. If substitutability or complementarity among items of a given class (e.g., women's slacks) is present, the correlation matrix should illustrate it. There were only eight observations (W_1 through W_8) for each item so that large correlations could be expected by chance alone. Table 5 contains the intercorrelation matrix for 12 slacks; there does not seem to be any discernible pattern among the large correlations or their signs. Dependencies cannot be exploited if their occurrence and sign cannot be anticipated.

Table 5

Intercorrelation Matrix for 12 Slacks (n=8)

```
 1.0000
 0.2244   1.0000
-0.4076   0.0346   1.0000
 0.5717   0.1441  -0.0853   1.0000
-0.1874  -0.4590   0.1841  -0.2323   1.0000
-0.4044  -0.1387   0.5046  -0.2003  -0.4396   1.0000
 0.7996   0.1576  -0.7198   0.3132  -0.4485  -0.1758   1.0000
-0.8268  -0.4012   0.2935  -0.6657   0.6595  -0.0124  -0.8032   1.0000
-0.4440  -0.6647   0.5558  -0.4508   0.4393   0.4331  -0.5666   0.5585   1.0000
 0.3165   0.4622  -0.7909   0.2344  -0.2293  -0.6480   0.5116  -0.2959  -0.8919   1.0000
 0.3255   0.2484   0.3214   0.3411  -0.5074   0.5489   0.3137  -0.6546  -0.0125  -0.3183   1.0000
-0.7225  -0.7063   0.3952  -0.4072   0.5069   0.2751  -0.7355   0.7864   0.8452  -0.6222  -0.2781
```

Exploring the disaggregation issue, Table 6 presents the size and color demand breakdown for six sweaters after normalizing the figures to percentages of each item/color's total demand. Item is denoted by letter and color by number in the first column.

Table 6

Sweater Demands by Color and Size

	Sizes				
Item/Color	1	2	3	4	Total across Size
A1	.136	.271	.271	.322	118
A2	.120	.227	.240	.413	75
B3	.179	.316	.253	.253	95
B4	.167	.317	.250	.267	60
C5	.146	.214	.301	.340	103
C6	.140	.211	.316	.333	57
D	.125	.299	.304	.272	184
E	.085	.245	.340	.330	200
F7	.213	.337	.247	.202	89
F8	.067	.500	.333	.100	30
Mean	.1378	.2937	.2855	.2832	
Std. Dev.	.0406	.0812	.0358	.0816	

The mean size breakdown and its standard deviation are also contained in the table; the size distribution does tend to remain relatively stable among items, with item F, color 8, being the most extreme outlier. As between colors, indicated by numbers in the first column, there is no tendency for demands to be evenly divided between colors. Thus it seems doubtful that much is gained from aggregating item-color demands into item demands; rather, it is likely that a need for action is obscured since low demands for one color may mask high demands for another. If inventory and demand by item are assessed they may appear to be in balance when the inventory by color is seriously out of balance. It has frequently been observed in practice that there are oversales of one stock-keeping unit and markdowns of another within one item, thereby reinforcing our suspicion that aggregating across colors can hide inventory imbalance (relative to demand) across colors.

Items returned by the customer have so far been ignored. In the mail-order situation returns constitute a significant factor as they both indicate that recent sales overstated true demand and they increase planned inventory. Although average returns for the situation studied were 10%, individual items had return rates as high as 38%, and the bad cases cause the worst losses. We suggest a simple monitoring scheme be implemented which considers the number of pieces of an item "exposed to risk of return," recognizing the delays in the return process, and which creates an "expected returns" forecast or allowance which is updated as experience is gained through the season. The details of this suggested scheme are in [14]; since no detailed data were available, no results can be presented.

APPENDIX 1

CALCULATION TO FIT STUDENT'S T DISTRIBUTION TO SAMPLE DATA

The procedure uses the ratio of two estimates of the standard deviation to find an approximate fit and then exploits two values usually found in the tabled t-distribution to get an exact fit to each period. Define the terms as follows:

$W_{.25}$ = the order statistic closest to the 25th percentile.

$W_{.75}$ = " " " " " " 75th " .

s = the standard deviation of all the sample W_{ij} for period j,

j being the period being fitted.

n = degrees of freedom, the parameter we are seeking to estimate.

If the sample is small all the values between the 25th percentile and the 75th should be plotted on normal paper, a line fitted, and the intersections of this line with the .25 and .75 ordinates used for $W_{.25}$ and $W_{.75}$. Then if $x = \frac{s}{.7413\,(|W_{.75} - W_{.25}|)}$, $n = \frac{2x^2}{x^2 - 1}$; $n = \infty$ if $x = 1$ (implying a Normal curve), and is undefined if $x < 1$ (a Normal curve is better than any t-curve). This formula gives a value for n that is too large because even between the quartiles the slope of the t-curve is steeper than the Normal curve. The value of x is properly $\frac{s}{a(|W_{.75} - W_{.25}|)}$ where a is dependent on n and is $\frac{1}{2b_n}$ where b_n = the value found in the P = .50 column of the t-table for n degrees of freedom.

n	x(using .7413)	x(exact)
2	1.779	-
3	1.402	1.743
4	1.264	1.414
5	1.193	1.290
6	1.149	1.225
7	1.122	1.183
8	1.102	1.145
9	1.085	1.133

The value of s needed to construct this table is found by dividing by two the value found in the P = .05 column of a two-tailed table of Student's t-distribution; this slightly under-estimates the value of s. In the exact column is given the value of s found analytically, which is $\sqrt{\frac{n}{n-2}}$, which is infinite for n=2 (it cannot be infinite for a sample).

APPENDIX 2

Poisson Test

The table shows in Col (1) each potential cell value from 1 to 125, in Col (2) the frequency with which D_{ij} actually took this value, and in Col (3) the theoretical expected frequency calculated on the Poisson assumption using parameters calculated from the sample data.

(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
1	13	13.398617	44	7	6.908802	87	1	0.493715
2	21	23.512201	45	9	6.387405	88	0	0.446995
3	35	32.278381	46	3	5.888939	89	1	0.403594
4	37	38.649101	47	3	5.418077	90	0	0.363506
5	48	42.507547	48	2	4.978673	91	0	0.326678
6	44	44.255444	49	2	4.573804	92	0	0.293026
7	46	44.507034	50	3	4.205319	93	0	0.262437
8	37	43.842658	51	5	3.874035	94	0	0.234774
9	45	42.684460	52	8	3.579723	95	0	0.209865
10	41	41.286583	53	3	3.321196	96	0	0.187536
11	38	39.779001	54	4	3.096387	97	0	0.167581
12	39	38.217130	55	2	2.902658	98	0	0.149811
13	40	36.618088	56	3	2.736818	99	0	0.134019
14	31	34.986505	57	3	2.595399	100	2	0.119999
15	35	33.327311	58	4	2.474803	101	0	0.107557
16	33	31.652198	59	3	2.371373	102	0	0.096512
17	32	29.979331	60	1	2.281767	103	0	0.086698
18	34	28.330754	61	2	2.202752	104	0	0.077963
19	28	26.728581	62	4	2.131453	105	0	0.070174
20	28	25.192585	63	5	2.065371	106	0	0.063221
21	16	23.738240	64	2	2.002445	107	0	0.056998
22	19	22.376541	65	2	1.940936	108	0	0.051447
23	16	21.113623	66	3	1.879572	109	0	0.046489
24	22	19.951787	67	0	1.817421	110	1	0.042078
25	17	18.889433	68	2	1.753868	111	0	0.038193
26	11	17.921878	69	1	1.688546	112	1	0.034791
27	17	17.041782	70	0	1.621375	113	0	0.031863
28	13	16.239716	71	1	1.552394	114	0	0.029396
29	22	15.504949	72	2	1.481847	115	0	0.027365
30	20	14.826136	73	1	1.410025	116	0	0.025773
31	16	14.191793	74	1	1.337292	117	0	0.024602
32	12	13.591280	75	2	1.264135	118	0	0.023835
33	10	13.014614	76	2	1.190947	119	0	0.023459
34	17	12.453138	77	0	1.118195	120	0	0.023447
35	10	11.899864	78	0	1.046287	121	0	0.023774
36	10	11.349292	79	0	0.975637	122	0	0.024421
37	11	10.797573	80	0	0.906632	123	0	0.025349
38	12	10.242466	81	1	0.839612	124	0	0.026522
39	13	9.683484	82	2	0.774896	125	0	0.027905
40	11	9.121628	83	0	0.712734	126	2	1.823294
41	6	8.559242	84	1	0.653371	127	6	4.465119
42	7	7.999889	85	1	0.596994	128	0	0.014604
43	7	7.448092	86	0	0.543746	129	1134	1127.696983

The cells numbered above 125 have special meaning as follows:

126: all possible cell values of 126 or more: It may be seen that D_{ij} twice exceeded 125 and was expected to do so 1.82 x.

127: the result for cell value zero is recorded here.

128, 129: these serve for bookkeeping and should be ignored.

REFERENCES

1. Aitchison, J. and Brown, J.A.C., The Lognormal Distribution, Cambridge University Press, Cambridge, England, 1967.

2. Chang, S.H. and Fyffe, D.E., "Estimation of Forecast Errors for Seasonal-Style-Goods Sales," Management Science, Vol. 18, No. 2 (October 1971), pp. B-89-B-96.

3. Crowston, W.B., Hausman, W.H., and Kampe, W.R., "Multistage Production for Stochastic Seasonal Demand," Sloan School of Management Working Paper #587-72 (April 1972).

4. Greenwood, M. and Yule, G.U., "An Inquiry Into the Nature of Frequency Distributions Representative of Multiple Happenings," Journal of Royal Statistical Society, 1920.

5. Hausman, W.H., "Sequential Decision Problems: A Model to Exploit Existing Forecasters," Management Science, Vol. 16, No. 2 (October 1969), pp. 93-111.

6. ____________ and Peterson, R., "Multiproduct Production Scheduling for Style Goods with Limited Capacity, Forecast Revisions and Terminal Delivery," Management Science (March 1972), pp. 370-383.

7. Hertz, D.B. and Shaffir, K.H., "A Forecasting Method for Management of Seasonal Style-Goods Inventories," Operations Research, Vol. 8 (1960), pp. 45-52.

8. Liff, Julius, "Inventory Management of Seasonal Style-Goods at the Retail Level," paper presented at the 36th National Meeting of the Operations Research Society of America, Miami Beach, Florida, November 1969.

9. Murray, G.R., Jr. and Silver, E.A., "A Bayesian Analysis of the Style Goods Inventory Problem," Management Science, Vol. 12, No. 11 (July 1966), pp. 785-797.

10. Praetz, P.D., "The Distribution of Share Price Changes," Journal of Business (January 1972), pp. 49-55.

11. Raiffa, H. and Schlaifer, R., Applied Statistical Decision Theory, Harvard University Press, 1961.

12. Ravindran, A., "Management of Seasonal Style-Goods Inventories," Operations Research, Vol. 20, No. 2 (March-April 1972), pp. 265-275.

13. Samuelson, P.A., "Proof that Properly Anticipated Prices Fluctuate Randomly," Industrial Management Review (Spring 1965), pp. 41-49.

14. Sides, R. St. G., An Inventory Model for Style Goods, Master's Thesis, Sloan School of Management, M.I.T., Cambridge, Mass., 1972.

15. Tukey, J. W., Exploratory Data Analysis, Addison-Wesley, 1970.

16. Wolfe, H. B., "A Model for Control of Style Merchandise", Industrial Management Review, Vol. 9, No. 2 (Winter 1968), pp. 69-82.

BASEMENT

Date Due

~~NOV [illegible]~~
MAY 21 76
FEB 20 79
MAR 22 '79
~~[illegible] '79~~
~~JAN [illegible] '79~~
AUG 13 '81
NOV 8 [illegible]
[illegible] 25 85
JUN 3 '86
SEP [illegible]
NO 27 '87

MAY 20 1990
AUG 10 1992
[illegible] 18 [illegible]
[illegible] 1999
FEB [illegible]
MAY 16 2000
NOV 2 [illegible] 2000
SEP [illegible]

Lib-26-67

BASEMENT

MIT LIBRARIES
3 9080 003 701 866

611-72

MIT LIBRARIES
3 9080 003 701 833

612-72

MIT LIBRARIES
3 9080 003 670 814

613-72

T-J5 143 w no.614- 72
Welsch, Roy El/A modification of the N
636673 D*BKS 00027764
3 9080 000 747 904

614

MIT LIBRARIES
3 9080 003 670 863

615-72

T-J5 143 w no.616- 72
Welsch, Roy El/The variances of regres
636653 D*BKS 00027762
3 9080 000 747 854

MIT LIBRARIES
3 9080 003 701 825

617-72

MIT LIBRARIES
3 9080 003 701 874

618-72

MIT LIBRARIES
3 9080 003 701 783

621-72

MIT LIBRARIES
3 9080 003 701 924

622-72

HD28.M414 no.623 72
Merton, Robert/Fallacy of the log-norm
636656 D*BKS 00027760
3 9080 000 747 789

623

HD28.M414 no.619-72
Plovnick, Mark/Expanding professional
636663 D*BKS 00027761
3 9080 000 747 813

[620 never issued]